

Distributed Delayed Stochastic Optimization

Alekh Agarwal John C. Duchi
 {alekh,jduchi}@eecs.berkeley.edu

Department of Electrical Engineering and Computer Sciences
 University of California, Berkeley

Abstract

We analyze the convergence of gradient-based optimization algorithms that base their updates on delayed stochastic gradient information. The main application of our results is to the development of gradient-based distributed optimization algorithms where a master node performs parameter updates while worker nodes compute stochastic gradients based on local information in parallel, which may give rise to delays due to asynchrony. We take motivation from statistical problems where the size of the data is so large that it cannot fit on one computer; with the advent of huge datasets in biology, astronomy, and the internet, such problems are now common. Our main contribution is to show that for smooth stochastic problems, the delays are asymptotically negligible and we can achieve order-optimal convergence results. In application to distributed optimization, we develop procedures that overcome communication bottlenecks and synchronization requirements. We show n -node architectures whose optimization error in stochastic problems—in spite of asynchronous delays—scales asymptotically as $\mathcal{O}(1/\sqrt{nT})$ after T iterations. This rate is known to be optimal for a distributed system with n nodes even in the absence of delays. We additionally complement our theoretical results with numerical experiments on a statistical machine learning task.

1 Introduction

We focus on stochastic convex optimization problems of the form

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad f(x) \quad \text{for} \quad f(x) := \mathbb{E}_P[F(x; \xi)] = \int_{\Xi} F(x; \xi) dP(\xi), \quad (1)$$

where $\mathcal{X} \subseteq \mathbb{R}^d$ is a closed convex set, P is a probability distribution over Ξ , and $F(\cdot; \xi)$ is convex for all $\xi \in \Xi$, so that f is convex. The goal is to find a parameter x that approximately minimizes f over $x \in \mathcal{X}$. Classical stochastic gradient algorithms [RM51, Pol87] iteratively update a parameter $x(t) \in \mathcal{X}$ by sampling $\xi \sim P$, computing $g(t) = \nabla F(x(t); \xi)$, and performing the update $x(t+1) = \Pi_{\mathcal{X}}(x(t) - \alpha(t)g(t))$, where $\Pi_{\mathcal{X}}$ denotes projection onto the set \mathcal{X} . In this paper, we analyze asynchronous gradient methods, where instead of receiving current information $g(t)$, the procedure receives out of date gradients $g(t-\tau(t)) = \nabla F(x(t-\tau(t)), \xi)$, where $\tau(t)$ is the (potentially random) delay at time t . The central contribution of this paper is to develop algorithms that—under natural assumptions about the functions F in the objective (1)—achieve asymptotically optimal rates for stochastic convex optimization in spite of delays.

Our model of delayed gradient information is particularly relevant in distributed optimization scenarios, where a master maintains the parameters x while workers compute stochastic gradients of the objective (1). The architectural assumption of a master with several worker nodes is natural for distributed computation, and other researchers have considered models similar to those in this paper [NBB01, LSZ09]. By allowing delayed and asynchronous updates, we can avoid synchronization issues that commonly handicap distributed systems.

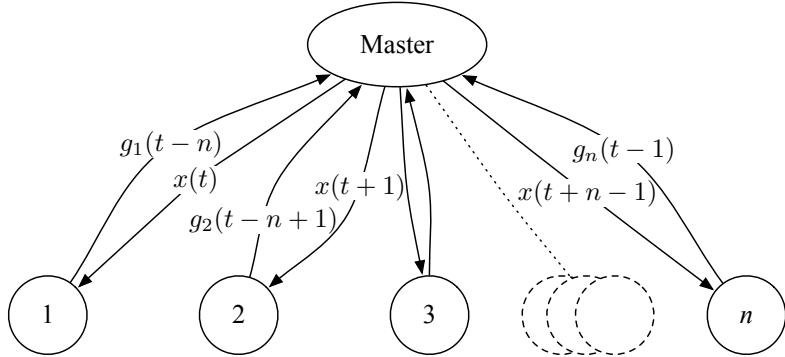


Figure 1. Cyclic delayed update architecture. Workers compute gradients cyclically and in parallel, passing out-of-date information to master. Master responds with current parameters. Diagram shows parameters and gradients communicated between rounds t and $t + n - 1$.

Certainly distributed optimization has been studied for several decades, tracing back at least to seminal work of Tsitsiklis and colleagues ([Tsi84, BT89]) on minimization of smooth functions where the parameter vector is distributed. More recent work has studied problems in which each processor or node i in a network has a local function f_i , and the goal is to minimize the sum $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$ [NO09, RNV10, JRJ09, DAW10]. Most prior work assumes as a *constraint* that data lies on several different nodes throughout a network. However, as Dekel et al. [DGBSX10a] first noted, in distributed stochastic settings independent realizations of a stochastic gradient can be computed concurrently, and it is thus possible to obtain an aggregated gradient estimate with lower variance. Using modern stochastic optimization algorithms (e.g. [JNT08, Lan10]), Dekel et al. give a series of reductions to show that in an n -node network it is possible to achieve a speedup of $\mathcal{O}(n)$ over a single-processor so long as the objective f is smooth.

Our work is closest to Nedić et al.’s asynchronous subgradient method [NBB01], which is an incremental gradient procedure in which gradient projection steps are taken using out-of-date gradients. See Figure 1 for an illustration. The asynchronous subgradient method performs non-smooth minimization and suffers an asymptotic penalty in convergence rate due to the delays: if the gradients are computed with a delay of τ , then the convergence rate of the procedure is $\mathcal{O}(\sqrt{\tau/T})$. The setting of distributed optimization provides an elegant illustration of the role played by the delay in convergence rates. As in Fig. 1, the delay τ can essentially be of order n in Nedić et al.’s setting, which gives a convergence rate of $\mathcal{O}(\sqrt{n/T})$. A simple centralized stochastic gradient algorithm attains a rate of $\mathcal{O}(1/\sqrt{T})$, which suggests something is amiss in the distributed algorithm. Langford et al. [LSZ09] rediscovered Nedić et al.’s results and attempted to remove the asymptotic penalty by considering smooth objective functions, though their approach has a technical error (see Appendix C), and even so they do not demonstrate any provable benefits of distributed computation. We analyze similar asynchronous algorithms, but we show that for smooth stochastic problems the delay is asymptotically negligible—the time τ does not matter—and in fact, with parallelization, delayed updates can give provable performance benefits.

We build on results of Dekel et al. [DGBSX10a], who show that when the objective f has Lipschitz-continuous gradients, then when n processors compute stochastic gradients in parallel using a common parameter x it is possible to achieve convergence rate $\mathcal{O}(1/\sqrt{Tn})$ so long as the processors are synchronized (under appropriate synchrony conditions, this holds nearly indepen-

dently of network topology). A variant of their approach is asymptotically robust to asynchrony so long as most processors remain synchronized for most of the time [DGBSX10b]. We show results similar to their initial discovery, but we analyze the effects of asynchronous gradient updates where all the nodes in the network can suffer delays. Application of our main results to the distributed setting provides convergence rates in terms of the number of nodes n in the network and the stochastic process governing the delays. Concretely, we show that under different assumptions on the network and delay process, we achieve convergence rates ranging from $\mathcal{O}(n^3/T + 1/\sqrt{Tn})$ to $\mathcal{O}(n/T + 1/\sqrt{Tn})$, which is $\mathcal{O}(1/\sqrt{nT})$ asymptotically in T . For problems with large n , we demonstrate faster rates ranging from $\mathcal{O}((n/T)^{2/3} + 1/\sqrt{Tn})$ to $\mathcal{O}(1/T^{2/3} + 1/\sqrt{Tn})$. In either case, the time necessary to achieve ϵ -optimal solution to the problem (1) is asymptotically $\mathcal{O}(1/n\epsilon^2)$, a factor of n —the size of the network—better than a centralized procedure in spite of delay.

The remainder of the paper is organized as follows. We begin by reviewing known algorithms for solving the stochastic optimization problem (1) and stating our main assumptions. Then in Section 3 we give abstract descriptions of our algorithms and state our main theoretical results, which we make concrete in Section 4 by formally placing the analysis in the setting of distributed stochastic optimization. We complement the theory in Section 5 with experiments on a real-world dataset, and proofs follow in the remaining sections.

Notation For the reader’s convenience, we collect our (mostly standard) notation here. We denote general norms by $\|\cdot\|$, and the dual norm $\|\cdot\|_*$ to the norm $\|\cdot\|$ is defined as $\|z\|_* := \sup_{x: \|x\| \leq 1} \langle z, x \rangle$. The subdifferential set of a function f is

$$\partial f(x) := \left\{ g \in \mathbb{R}^d \mid f(y) \geq f(x) + \langle g, y - x \rangle \text{ for all } y \in \text{dom } f \right\}$$

We use the shorthand $\|\partial f(x)\|_* := \sup_{g \in \partial f(x)} \|g\|_*$. A function f is G -Lipschitz with respect to the norm $\|\cdot\|$ on \mathcal{X} if for all $x, y \in \mathcal{X}$, $|f(x) - f(y)| \leq G \|x - y\|$. For convex f , this is equivalent to $\|\partial f(x)\|_* \leq G$ for all $x \in \mathcal{X}$ (e.g. [HUL96a]). A function f is L -smooth on \mathcal{X} if ∇f is Lipschitz continuous with respect to the norm $\|\cdot\|$, defined as

$$\|\nabla f(x) - \nabla f(y)\|_* \leq L \|x - y\|, \quad \text{equivalently,} \quad f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{L}{2} \|x - y\|^2.$$

For convex differentiable h , the Bregman divergence [Bre67] between x and y is defined as

$$D_h(x, y) := h(x) - h(y) - \langle \nabla h(y), x - y \rangle. \quad (2)$$

A convex function h is c -strongly convex with respect to a norm $\|\cdot\|$ over \mathcal{X} if

$$h(y) \geq h(x) + \langle g, y - x \rangle + \frac{c}{2} \|x - y\|^2 \quad \text{for all } x, y \in \mathcal{X} \text{ and } g \in \partial h(x). \quad (3)$$

We use $[n]$ to denote the set of integers $\{1, \dots, n\}$.

2 Setup and Algorithms

In this section we set up and recall the delay-free algorithms underlying our approach. We then give the appropriate delayed versions of these algorithms, which we analyze in the sequel.

2.1 Setup and Delay-free Algorithms

To build intuition for the algorithms we analyze, we first describe two closely related first-order algorithms: the dual averaging algorithm of Nesterov [Nes09] and the mirror descent algorithm of Nemirovski and Yudin [NY83], which is analyzed further by Beck and Teboulle [BT03]. We begin by collecting notation and giving useful definitions. Both algorithms are based on a proximal function $\psi(x)$, where it is no loss of generality to assume that $\psi(x) \geq 0$ for all $x \in \mathcal{X}$. We assume ψ is 1-strongly convex (by scaling, this is no loss of generality). By definitions (2) and (3), the divergence D_ψ satisfies $D_\psi(x, y) \geq \frac{1}{2} \|x - y\|^2$.

In the oracle model of stochastic optimization that we assume, at time t both algorithms query an oracle at the point $x(t)$, and the oracle then samples $\xi(t)$ i.i.d. from the distribution P and returns $g(t) \in \partial F(x(t); \xi(t))$. The dual averaging algorithm [Nes09] updates a dual vector $z(t)$ and primal vector $x(t) \in \mathcal{X}$ via

$$z(t+1) = z(t) + g(t) \quad \text{and} \quad x(t+1) = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle z(t+1), x \rangle + \frac{1}{\alpha(t+1)} \psi(x) \right\}, \quad (4)$$

while mirror descent [NY83, BT03] performs the update

$$x(t+1) = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle g(t), x \rangle + \frac{1}{\alpha(t)} D_\psi(x, x(t)) \right\}. \quad (5)$$

Both make a linear approximation to the function being minimized—a global approximation in the case of the dual averaging update (4) and a more local approximation for mirror descent (5)—while using the proximal function ψ to regularize the points $x(t)$.

We now state the two essentially standard assumptions [JNT08, Lan10, Xia10] we most often make about the stochastic optimization problem (1), after which we recall the convergence rates of the algorithms (4) and (5).

Assumption A (Lipschitz Functions). *For P -a.e. ξ , the function $F(\cdot; \xi)$ is convex. Moreover, for any $x \in \mathcal{X}$, $\mathbb{E}[\|\partial F(x; \xi)\|_*^2] \leq G^2$.*

In particular, Assumption A implies that f is G -Lipschitz continuous with respect to the norm $\|\cdot\|$ and that f is convex. Our second assumption has been used to show rates of convergence based on the variance of a gradient estimator for stochastic optimization problems (e.g. [JNT08, Lan10]).

Assumption B (Smooth Functions). *The function f defined in (1) has L -Lipschitz continuous gradient, and for all $x \in \mathcal{X}$ the variance bound $\mathbb{E}[\|\nabla f(x) - \nabla F(x; \xi)\|_*^2] \leq \sigma^2$ holds.¹*

Several commonly used functions satisfy the above assumptions, for example:

- (i) The *logistic loss*: $F(x; \xi) = \log[1 + \exp(\langle x, \xi \rangle)]$, the objective for logistic regression in statistics (e.g. [HTF01]). The objective F satisfies Assumptions A and B so long as $\|\xi\|$ is bounded.
- (ii) *Least squares or linear regression*: $F(x; \xi) = (a - \langle x, b \rangle)^2$ where $\xi = (a, b)$ for $a \in \mathbb{R}^d$ and $b \in \mathbb{R}$, satisfies Assumptions A and B as long as ξ is bounded and \mathcal{X} is compact.

We also make a standard compactness assumption on the optimization set \mathcal{X} .

¹If f is differentiable, then $F(\cdot; \xi)$ is differentiable for P -a.e. ξ , and conversely, but F need not be smoothly differentiable [Ber73]. Since $\nabla F(x; \xi)$ exists for P -a.e. ξ , we will write $\nabla F(x; \xi)$ with no loss of generality.

Assumption C (Compactness). *For $x^* \in \operatorname{argmin}_{x \in \mathcal{X}} f(x)$ and $x \in \mathcal{X}$, the bounds $\psi(x^*) \leq R^2/2$ and $D_\psi(x^*, x) \leq R^2$ both hold.*

Under Assumptions A or B in addition to Assumption C, the updates (4) and (5) have known convergence rates. Define the time averaged vector $\hat{x}(T)$ as

$$\hat{x}(T) := \frac{1}{T} \sum_{t=1}^T x(t+1). \quad (6)$$

Then under Assumption A, both algorithms satisfy

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}\left(\frac{RG}{\sqrt{T}}\right) \quad (7)$$

for the stepsize choice $\alpha(t) = R/(G\sqrt{t})$ (e.g. [Nes09, Xia10, NJLS09]). The result (7) is sharp to constant factors in general [NY83, ABRW10], but can be further improved under Assumption B. Building on work of Juditsky et al. [JNT08] and Lan [Lan10], Dekel et al. [DGBSX10a, Appendix A] show that under Assumptions B and C the stepsize choice $\alpha(t)^{-1} = L + \eta(t)$, where $\eta(t)$ is a damping factor set to $\eta(t) = \sigma R\sqrt{t}$, yields for either of the updates (4) or (5) the convergence rate

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}\left(\frac{LR^2}{T} + \frac{\sigma R}{\sqrt{T}}\right). \quad (8)$$

2.2 Delayed Optimization Algorithms

We now turn to extending the dual averaging (4) and mirror descent (5) updates to the setting in which instead descent (5) updates to the setting in which instead of receiving a current gradient $g(t)$ at time t , the procedure receives a gradient $g(t - \tau(t))$, that is, a stochastic gradient of the objective (1) computed at the point $x(t - \tau(t))$. In the simplest case, the delays are uniform and $\tau(t) \equiv \tau$ for all t , but in general the delays may be a non-i.i.d. stochastic process. Our analysis admits any sequence $\tau(t)$ of delays as long as the mapping $t \mapsto \tau(t)$ satisfies $\mathbb{E}[\tau(t)] \leq B < \infty$. We also require that each update happens once, i.e., $t \mapsto t - \tau(t)$ is one-to-one, though this second assumption is easily satisfied.

Recall that the problems we consider are stochastic optimization problems of the form (1). Under the assumptions above, we extend the mirror descent and dual averaging algorithms in the simplest way: we replace $g(t)$ with $g(t - \tau(t))$. For dual averaging (c.f. the update (4)) this yields

$$z(t+1) = z(t) + g(t - \tau(t)) \quad \text{and} \quad x(t+1) = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle z(t+1), x \rangle + \frac{1}{\alpha(t+1)} \psi(x) \right\}, \quad (9)$$

while for mirror descent (c.f. the update (5)) we have

$$x(t+1) = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle g(t - \tau(t)), x \rangle + \frac{1}{\alpha(t)} D_\psi(x, x(t)) \right\}. \quad (10)$$

A generalization of Nedić et al.'s results [NBB01] by combining their techniques with the convergence proofs of dual averaging [Nes09] and mirror descent [BT03] is as follows. Under Assumptions A and C, so long as $\mathbb{E}[\tau(t)] \leq B < \infty$ for all t , choosing $\alpha(t) = \frac{R}{G\sqrt{Bt}}$ gives rate

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}\left(\frac{RG\sqrt{B}}{\sqrt{T}}\right). \quad (11)$$

3 Convergence rates for delayed optimization of smooth functions

In this section, we state and discuss several results for asynchronous stochastic gradient methods. We give two sets of theorems. The first are for the asynchronous method when we make updates to the parameter vector x using one stochastic subgradient, according to the update rules (9) or (10). The second method involves using several stochastic subgradients for every update, each with a potentially different delay, which gives sharper results that we present in Section 3.2.

3.1 Simple delayed optimization

Intuitively, the \sqrt{B} -penalty due to delays for non-smooth optimization arises from the fact that subgradients can change drastically when measured at slightly different locations, so a small delay can introduce significant inaccuracy. To overcome the delay penalty, we now turn to the smoothness assumption B as well as the Lipschitz condition A (we assume both of these conditions along with Assumption C hold for all the theorems). In the smooth case, delays mean that stale gradients are only slightly perturbed, since our stochastic algorithms constrain the variability of the points $x(t)$. As we show in the proofs of the remaining results, the error from delay essentially becomes a second order term: the penalty is asymptotically negligible. We study both update rules (9) and (10), and we set $\alpha(t) = \frac{1}{L+\eta(t)}$. Here $\eta(t)$ will be chosen to both control the effects of delays and for errors from stochastic gradient information. We prove the following theorem in Sec. 6.1.

Theorem 1. *Let the sequence $x(t)$ be defined by the update (9). Define the stepsize $\eta(t) \propto \sqrt{t+\tau}$ or let $\eta(t) \equiv \eta$ for all t . Then*

$$\mathbb{E} \left[\sum_{t=1}^T f(x(t+1)) \right] - T f(x^*) \leq \frac{1}{\alpha(T+1)} R^2 + \frac{\sigma^2}{2} \sum_{t=1}^T \frac{1}{\eta(t)} + 2LG^2(\tau+1)^2 \sum_{t=1}^T \frac{1}{\eta(t-\tau)^2} + 2\tau GR.$$

The mirror descent update (10) exhibits similar convergence properties, and we prove the next theorem in Sec. 6.2.

Theorem 2. *Use the conditions of Theorem 1 but generate $x(t)$ by the update (10). Then*

$$\mathbb{E} \left[\sum_{t=1}^T f(x(t+1)) \right] - T f(x^*) \leq 2LR^2 + R^2[\eta(1) + \eta(T)] + \frac{\sigma^2}{2} \sum_{t=1}^T \frac{1}{\eta(t)} + 2LG^2(\tau+1)^2 \sum_{t=1}^T \frac{1}{\eta(t-\tau)^2} + 2\tau GR$$

In each of the above theorems, we can set $\eta(t) = \sigma\sqrt{t+\tau}/R$. As immediate corollaries, we recall the definition (6) of the averaged sequence of $x(t)$ and use convexity to see that

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O} \left(\frac{LR^2 + \tau GR}{T} + \frac{\sigma R}{\sqrt{T}} + \frac{LG^2 \tau^2 R^2 \log T}{\sigma^2 T} \right)$$

for either update rule. In addition, we can allow the delay $\tau(t)$ to be random:

Corollary 1. *Let the conditions of Theorem 1 or 2 hold, but allow $\tau(t)$ to be a random mapping such that $\mathbb{E}[\tau(t)^2] \leq B^2$ for all t . With the choice $\eta(t) = \sigma\sqrt{T}/R$ the updates (9) and (10) satisfy*

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O} \left(\frac{LR^2 + B^2 GR}{T} + \frac{\sigma R}{\sqrt{T}} + \frac{LG^2 B^2 R^2}{\sigma^2 T} \right).$$

We provide the proof of the corollary in Sec. 6.3. The take-home message from the above corollaries, as well as Theorems 1 and 2, is that the penalty in convergence rate due to the delay $\tau(t)$ is asymptotically negligible. As we discuss in greater depth in the next section, this has favorable implications for robust distributed stochastic optimization algorithms.

3.2 Combinations of delays

In some scenarios—including distributed settings similar to those we discuss in the next section—the procedure has access not to only a single delayed gradient but to several with different delays. To abstract away the essential parts of this situation, we assume that the procedure receives n gradients g_1, \dots, g_n , where each has a potentially different delay $\tau(i)$. Now let $\lambda = (\lambda_i)_{i=1}^n$ belong to the probability simplex, though we leave λ 's values unspecified for now. Then the procedure performs the following updates at time t : for dual averaging,

$$z(t+1) = z(t) + \sum_{i=1}^n \lambda_i g_i(t - \tau(i)) \quad \text{and} \quad x(t+1) = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle z(t+1), x \rangle + \frac{1}{\alpha(t+1)} \psi(x) \right\} \quad (12)$$

while for mirror descent, the update is

$$g_\lambda(t) = \sum_{i=1}^n \lambda_i g_i(t - \tau(i)) \quad \text{and} \quad x(t+1) = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \langle g_\lambda(t), x \rangle + \frac{1}{\alpha(t)} D_\psi(x, x(t)) \right\}. \quad (13)$$

The next two theorems build on the proofs of Theorems 1 and 2, combining several techniques. We provide the proof of Theorem 3 in Sec. 7, omitting the proof of Theorem 4 as it follows in a similar way from Theorem 2.

Theorem 3. *Let the sequence $x(t)$ be defined by the update (12). Under assumptions A, B and C, let $\frac{1}{\alpha(t)} = L + \eta(t)$ and $\eta(t) \propto \sqrt{t + \tau}$ or $\eta(t) \equiv \eta$ for all t . Then*

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T f(x(t+1)) - T f(x^*) \right] &\leq LR^2 + \eta(T)R^2 + 2 \sum_{i=1}^n \lambda_i \tau(i) GR + 2 \sum_{i=1}^n \lambda_i L G^2 (\tau(i) + 1)^2 \sum_{t=1}^T \frac{1}{\eta(t - \tau)^2} \\ &\quad + \sum_{t=1}^T \frac{1}{2\eta(t)} \mathbb{E} \left\| \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))] \right\|_*^2. \end{aligned}$$

Theorem 4. *Use the same conditions as Theorem 3, but assume that $x(t)$ is defined by the update (13) and $D_\psi(x^*, x) \leq R^2$ for all $x \in \mathcal{X}$. Then*

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T f(x(t+1)) - T f(x^*) \right] &\leq 2R^2(L + \eta(T)) + 2 \sum_{i=1}^n \lambda_i \tau(i) GR + 2 \sum_{i=1}^n \lambda_i L G^2 (\tau(i) + 1)^2 \sum_{t=1}^T \frac{1}{\eta(t - \tau)^2} \\ &\quad + \sum_{t=1}^T \frac{1}{2\eta(t)} \mathbb{E} \left\| \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))] \right\|_*^2. \end{aligned}$$

The consequences of Theorems 3 and 4 are powerful, as we illustrate in the next section.

4 Distributed Optimization

We now turn to what we see as the main purpose and application of the above results: developing robust and efficient algorithms for distributed stochastic optimization. Our main motivations here are machine learning and statistical applications where the data is so large that it cannot fit on a single computer. Examples of the form (1) include logistic regression (for background, see [HTF01]), where the task is to learn a linear classifier that assigns labels in $\{-1, +1\}$ to a series of examples, in which case we have the objective $F(x; \xi) = \log[1 + \exp(\langle \xi, x \rangle)]$ as described in Sec. 2.1(i); or linear regression, where $\xi = (a, b) \in \mathbb{R}^d \times \mathbb{R}$ and $F(x; \xi) = \frac{1}{2}[b - \langle a, x \rangle]^2$ as described in Sec. 2.1(ii). Both objectives satisfy assumptions A and B as discussed earlier. We consider both stochastic and online/streaming scenarios for such problems. In the simplest setting, the distribution P in the objective (1) is the empirical distribution over an observed dataset, that is,

$$f(x) = \frac{1}{N} \sum_{i=1}^N F(x; \xi_i).$$

We divide the N samples among n workers so that each worker has an N/n -sized subset of data. In streaming applications, the distribution P is the unknown distribution generating the data, and each worker receives a stream of independent data points $\xi \sim P$. Worker i uses its subset of the data, or its stream, to compute g_i , an estimate of the gradient ∇f of the global f . We make the simplifying assumption that g_i is an unbiased estimate of $\nabla f(x)$, which is satisfied, for example, when each worker receives an independent stream of samples or computes the gradient g_i based on samples picked at random without replacement from its subset of the data.

The architectural assumptions we make are natural and based off of master/worker topologies, but the convergence results in Section 3 allow us to give procedures robust to delay and asynchrony. We consider two protocols: in the first, workers compute and communicate asynchronously and independently with the master, and in the second, workers are at different distances from the master and communicate with time lags proportional to their distances. We show in the latter part of this section that the convergence rates of each protocol, when applied in an n -node network, are $\mathcal{O}(1/\sqrt{nT})$ for n -node networks (though lower order terms are different for each).

Before describing our architectures, we note that perhaps the simplest master-worker scheme is to have each worker simultaneously compute a stochastic gradient and send it to the master, which takes a gradient step on the averaged gradient. While the n gradients are computed in parallel, accumulating and averaging n gradients at the master takes $\Omega(n)$ time, offsetting the gains of parallelization. Thus we consider alternate architectures that are robust to delay.

Cyclic Delayed Architecture This protocol is the delayed update algorithm mentioned in the introduction, and it parallelizes computation of (estimates of) the gradient $\nabla f(x)$. Formally, worker i has parameter $x(t)$ and computes $g_i(t) = F(x(t); \xi_i(t))$, where $\xi_i(t)$ is a random variable sampled at worker i from the distribution P . The master maintains a parameter vector $x \in \mathcal{X}$. The algorithm proceeds in rounds, cyclically pipelining updates. The algorithm begins by initiating gradient computations at different workers at slightly offset times. At time t , the master receives gradient information at a τ -step delay from some worker, performs a parameter update, and passes the updated central parameter $x(t+1)$ back to the worker. Other workers do not see this update and continue their gradient computations on stale parameter vectors. In the simplest case, each

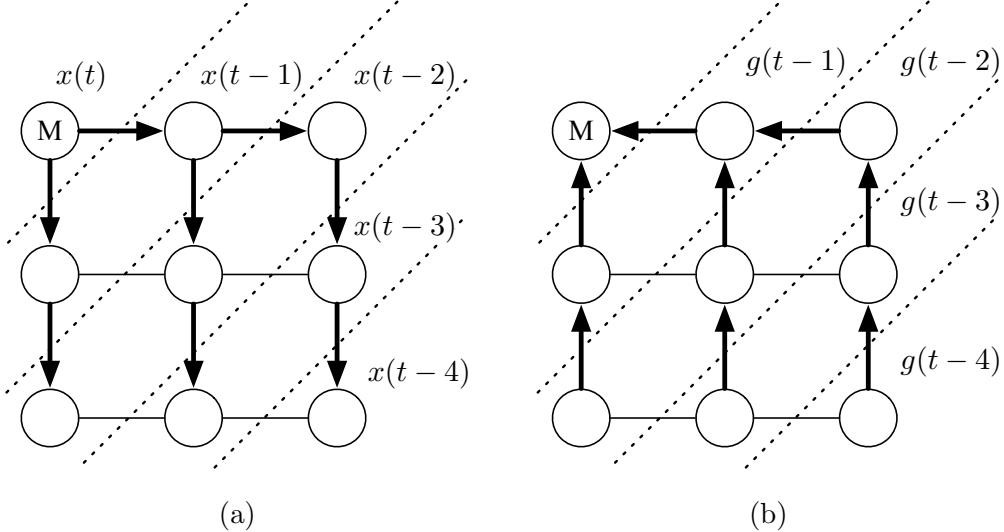


Figure 2. Master-worker averaging network. (a): parameters stored at different distances from master node at time t . A node at distance d from master has the parameter $x(t-d)$. (b): gradients computed at different nodes. A node at distance d from master computes gradient $g(t-d)$.

node suffers a delay of $\tau = n$, though our earlier analysis applies to random delays throughout the network as well. Recall Fig. 1 for a graphic description of the process.

Locally Averaged Delayed Architecture At a high level, the protocol we now describe combines the delayed updates of the cyclic delayed architecture with averaging techniques of previous work [NO09, DAW10]. We assume a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a set of n nodes (workers) and \mathcal{E} are the edges between the nodes. We select one of the nodes as the master, which maintains the parameter vector $x(t) \in \mathcal{X}$ over time.

The algorithm works via a series of multicasting and aggregation steps on a spanning tree rooted at the master node. In the first phase, the algorithm broadcasts from the root towards the leaves. At step t the master sends its current parameter vector $x(t)$ to its immediate neighbors. Simultaneously, every other node broadcasts its current parameter vector (which, for a depth d node, is $x(t-d)$) to its children in the spanning tree. See Fig. 2(a). Every worker receives its new parameter and computes its local gradient at this parameter. The second part of the communication in a given iteration proceeds from leaves toward the root. The leaf nodes communicate their gradients to their parents. The parent takes the gradients of the leaf nodes from the previous round (received at iteration $t-1$) and averages them with its own gradient, passing this averaged gradient back up the tree. Again simultaneously, each node takes the averaged gradient vectors of its children from the previous rounds, averages them with its current gradient vector, and passes the result up the spanning tree. See Fig. 2(b) and Fig. 3 for a visual description.

Slightly more formally, associated with each node $i \in \mathcal{V}$ is a delay $\tau(i)$, which is (generally) twice its distance from the master. Fix an iteration t . Each node $i \in \mathcal{V}$ has an out of date parameter vector $x(t - \tau(i)/2)$, which it sends further down the tree to its children. So, for example, the master node sends the vector $x(t)$ to its children, which send the parameter vector $x(t-1)$ to their

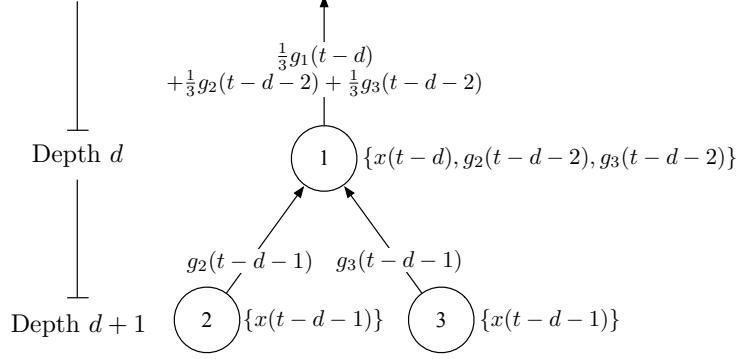


Figure 3. Communication of gradient information toward master node at time t from node 1 at distance d from master. Information stored at time t by node i in brackets to right of node.

children, which in turn send $x(t-2)$ to their children, and so on. Each node computes

$$g_i(t - \tau(i)/2) = \nabla F(x(t - \tau(i)/2); \xi_i(t)),$$

where $\xi_i(t)$ is a random variable sampled at node i from the distribution P . The communication back up the hierarchy proceeds as follows: the leaf nodes in the tree (say at depth d) send the gradient vectors $g_i(t-d)$ to their immediate parents in the tree. At the previous iteration $t-1$, the parent nodes received $g_i(t-d-1)$ from their children, which they average with their own gradients $g_i(t-d+1)$ and pass to their parents, and so on. The master node at the root of the tree receives an average of delayed gradients from the entire tree, with each gradient having a potentially different delay, giving rise to updates of the form (12) or (13).

4.1 Convergence rates for delayed distributed minimization

Having described our architectures, we can now give corollaries to the theoretical results from the previous sections that show it is possible to achieve asymptotically faster rates (over centralized procedures) using distributed algorithms even without imposing synchronization requirements. We allow workers to pipeline updates by computing asynchronously and in parallel, so each worker can compute low variance estimate of the gradient $\nabla f(x)$.

We begin with a simple corollary to the results in Sec. 3.1. We ignore the constants L , G , R , and σ , which are not dependent on the characteristics of the network. We also assume that each worker uses m independent samples of $\xi \sim P$ to compute the stochastic gradient as

$$g_i(t) = \frac{1}{m} \sum_{j=1}^m \nabla F(x(t); \xi_i(j)).$$

Using the cyclic protocol as in Fig. 1, Theorems 1 and 2 give the following result.

Corollary 2. *Let $\psi(x) = \frac{1}{2} \|x\|_2^2$, assume the conditions in Corollary 1, and assume that each worker uses m samples $\xi \sim P$ to compute the gradient it communicates to the master. Then with the choice $\eta(t) = \sqrt{T}/\sqrt{m}$ either of the updates (9) or (10) satisfy*

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}\left(\frac{B^2}{T} + \frac{1}{\sqrt{Tm}} + \frac{B^2m}{T}\right).$$

Proof The corollary follows straightforwardly from the realization that the variance $\sigma^2 = \mathbb{E}[\|\nabla f(x) - g_i(t)\|_2^2] = \mathbb{E}[\|\nabla f(x) - \nabla F(x; \xi)\|_2^2]/m = \mathcal{O}(1/m)$ when workers use m independent stochastic gradient samples. \square

In the above corollary, so long as the bound on the delay B satisfies, say, $B = o(T^{1/4})$, then the last term in the bound is asymptotically negligible, and we achieve a convergence rate of $\mathcal{O}(1/\sqrt{Tm})$.

The cyclic delayed architecture has the drawback that information from a worker can take $\mathcal{O}(n)$ time to reach the master. While the algorithm is robust to delay and does not need lock-step coordination of workers, the downside of the architecture is that the essentially n^2m/T term in the bounds above can be quite large. Indeed, if each worker computes its gradient over m samples with $m \approx n$ —say to avoid idling of workers—then the cyclic architecture has convergence rate $\mathcal{O}(n^3/T + 1/\sqrt{nT})$. For moderate T or large n , the delay penalty n^3/T may dominate $1/\sqrt{nT}$, offsetting the gains of parallelization.

To address the large n drawback, we turn our attention to the locally averaged architecture described by Figs. 2 and 3, where delays can be smaller since they depend only on the height of a spanning tree in the network. The algorithm requires more synchronization than the cyclic architecture but still performs limited local communication. Each worker computes $g_i(t - \tau(i)) = \nabla F(x(t - \tau(i)); \xi_i(t))$ where $\tau(i)$ is the delay of worker i from the master and $\xi_i \sim P$. As a result of the communication procedure, the master receives a convex combination of the stochastic gradients evaluated at each worker i , for which we gave results in Section 3.2.

In this architecture, the master receives gradients of the form $g_\lambda(t) = \sum_{i=1}^n \lambda_i g_i(t - \tau(i))$ for some λ in the simplex, which puts us in the setting of Theorems 3 and 4. We now make the reasonable assumption that the gradient errors $\nabla f(x(t)) - g_i(t)$ are uncorrelated across the nodes in the network.² In statistical applications, for example, each worker may own independent data or receive streaming data from independent sources; more generally, each worker can simply receive independent samples $\xi_i \sim P$. We also set $\psi(x) = \frac{1}{2} \|x\|_2^2$, and observe

$$\mathbb{E} \left\| \sum_{i=1}^n \lambda_i \nabla f(x(t - \tau(i))) - g_i(t - \tau(i)) \right\|_2^2 = \sum_{i=1}^n \lambda_i^2 \mathbb{E} \|\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))\|_2^2.$$

This gives the following corollary to Theorems 3 and 4.

Corollary 3. Set $\lambda_i = \frac{1}{n}$ for all i , $\psi(x) = \frac{1}{2} \|x\|_2^2$, and $\eta(t) = \sigma\sqrt{t + \tau}/R\sqrt{n}$. Let $\bar{\tau}$ and $\bar{\tau}^2$ denote the average of the delays $\tau(i)$ and $\tau(i)^2$, respectively. Under the conditions of Theorem 3 or 4,

$$\mathbb{E} \left[\sum_{t=1}^T f(x(t+1)) - T f(x^*) \right] = \mathcal{O} \left(LR^2 + \bar{\tau} GR + \frac{LG^2 R^2 n \bar{\tau}^2}{\sigma^2} \log T + \frac{R\sigma}{\sqrt{n}} \sqrt{T} \right).$$

The $\log T$ multiplier can be reduced to a constant if we set $\eta(t) \equiv \sigma\sqrt{T}/R\sqrt{n}$. By using the averaged sequence $\hat{x}(T)$ (6), Jensen’s inequality gives that asymptotically $\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}(1/\sqrt{Tn})$, which is an optimal dependence on the number of samples ξ calculated by the method. We also observe in this architecture, the delay τ is bounded by the graph diameter D , giving us the bound:

$$\mathbb{E} \left[\sum_{t=1}^T f(x(t+1)) - T f(x^*) \right] = \mathcal{O} \left(LR^2 + DGR + \frac{LG^2 R^2 n D^2}{\sigma^2} \log T + \frac{R\sigma}{\sqrt{n}} \sqrt{T} \right). \quad (14)$$

²Similar results continue to hold under weak correlation.

The above corollaries are general and hold irrespective of the relative costs of communication and computation. However, with knowledge of the costs, we can adapt the stepsizes slightly to give better rates of convergence when n is large or communication to the master node is expensive. For now, we focus on the cyclic architecture (the setting of Corollary 2), though the same principles apply to the local averaging scheme. Let C denote the cost of communicating between the master and workers in terms of the time to compute a single gradient sample, and assume that we set $m = Cn$, so that no worker node has idle time. For simplicity, we let the delay be non-random, so $B = \tau = n$. Consider the choice $\eta(t) = \eta\sqrt{T/(Cn)}$ for the damping stepsizes, where $\eta \geq 1$. This setting in Theorem 1 gives

$$\mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}\left(\frac{\eta^2 Cn^3}{T} + \frac{\eta}{\sqrt{TCn}} + \frac{1}{\eta\sqrt{TCn}}\right) = \mathcal{O}\left(\frac{\eta^2 Cn^3}{T} + \frac{\eta}{\sqrt{TCn}}\right),$$

where the last equality follows since $\eta \geq 1$. Optimizing for η on the right yields

$$\eta = \max\left\{\frac{n^{7/6}}{T^{1/6}C^{1/2}}, 1\right\} \quad \text{and} \quad \mathbb{E}[f(\hat{x}(T))] - f(x^*) = \mathcal{O}\left(\min\left\{\frac{n^{2/3}}{T^{2/3}}, \frac{n^3}{T}\right\} + \frac{1}{\sqrt{TCn}}\right). \quad (15)$$

The convergence rates thus follow two regimes. When $T \leq n^7/C^3$, we have convergence rate $\mathcal{O}(n^{2/3}/T^{2/3})$, while once $T > n^7/C^3$, we attain $\mathcal{O}(1/\sqrt{TCn})$ convergence. Roughly, in time proportional to TC , we achieve optimization error $1/\sqrt{TCn}$, which is order-optimal given that we can compute a total of TCn stochastic gradients [ABRW10]. The scaling of this bound is nicer than that previously: the dependence on network size is at worst $n^{2/3}$, which we obtain by increasing the damping factor $\eta(t)$ —and hence decreasing the stepsize $\alpha(t) = 1/(L + \eta(t))$ —relative to the setting of Corollary 2. We remark that applying the same technique to Corollary 3 gives convergence rate scaling as the smaller of $\mathcal{O}((D/T)^{2/3} + 1/\sqrt{TCn})$ and $\mathcal{O}((nCD/T + 1/\sqrt{TCn})$. Since the diameter $D \leq n$, this is faster than the cyclic architecture’s bound (15).

4.2 Running-time comparisons

Having derived the rates of convergence of the different distributed procedures above, we now explicitly study the running times of the centralized stochastic gradient algorithms (4) and (5), the cyclic delayed protocol with the updates (9) and (10), and the locally averaged architecture with the updates (12) and (13). To make comparisons more cleanly, we avoid constants, assuming without loss that the variance bound σ^2 on $\mathbb{E}\|\nabla f(x) - \nabla F(x; \xi)\|^2$ is 1, and that sampling $\xi \sim P$ and evaluating $\nabla F(x; \xi)$ requires one unit of time. Noting that $\mathbb{E}[\nabla F(x; \xi)] = \nabla f(x)$, it is clear that if we receive m uncorrelated samples of ξ , the variance $\mathbb{E}\|\nabla f(x) - \frac{1}{m} \sum_{j=1}^m \nabla F(x; \xi_j)\|_2^2 \leq \frac{1}{m}$.

Now we state our assumptions on the relative times used by each algorithm. Let T be the number of units of time allocated to each algorithm, and let the centralized, cyclic delayed and locally averaged delayed algorithms complete T_{cent} , T_{cycle} and T_{dist} iterations, respectively, in time T . It is clear that $T_{\text{cent}} = T$. We assume that the distributed methods use m_{cycle} and m_{dist} samples of $\xi \sim P$ to compute stochastic gradients and that the delay τ of the cyclic algorithm is n . For concreteness, we assume that communication is of the same order as computing the gradient of one sample $\nabla F(x; \xi)$ so that $C = 1$. In the cyclic setup of Sec. 3.1, it is reasonable to assume that $m_{\text{cycle}} = \Omega(n)$ to avoid idling of workers (Theorems 1 and 2, as well as the bound (15), show it is asymptotically beneficial to have m_{cycle} larger, since $\sigma_{\text{cycle}}^2 = 1/m_{\text{cycle}}$). For $m_{\text{cycle}} = \Omega(n)$, the master requires $\frac{m_{\text{cycle}}}{n}$ units of time to receive one gradient update, so $\frac{m_{\text{cycle}}}{n}T_{\text{cycle}} = T$. In

Centralized (4, 5)	$\mathbb{E}f(\hat{x}) - f(x^*) = \mathcal{O}\left(\sqrt{\frac{1}{T}}\right)$
Cyclic (9, 10)	$\mathcal{O}\left(\min\left(\frac{n^{2/3}}{T^{2/3}}, \frac{n^3}{T}\right) + \frac{1}{\sqrt{Tn}}\right)$
Local (12, 13)	$\mathbb{E}f(\hat{x}) - f(x^*) = \mathcal{O}\left(\min\left(\frac{D^{2/3}}{T^{2/3}}, \frac{n\bar{\tau}^2}{T}\right) + \frac{1}{\sqrt{nT}}\right)$

Table 1. Upper bounds on $\mathbb{E}f(\hat{x}) - f(x^*)$ for three computational architectures, where \hat{x} is the output of each algorithm after T units of time. Each algorithm runs for the amount of time it takes a centralized stochastic algorithm to perform T iterations as in (16). Here D is the diameter of the network, n is the number of nodes, and $\bar{\tau}^2 = \frac{1}{n} \sum_{i=1}^n \tau(i)^2$ is the average squared communication delay for the local averaging architecture. Bounds for the cyclic architecture assume delay $\tau = n$.

the locally delayed framework, if each node uses m_{dist} samples to compute a gradient, the master receives a gradient every m_{dist} units of time, and hence $m_{\text{dist}}T_{\text{dist}} = T$. Further, $\sigma_{\text{dist}}^2 = 1/m_{\text{dist}}$. We summarize our assumptions by saying that in T units of time, each algorithm performs the following number of iterations:

$$T_{\text{cent}} = T, \quad T_{\text{cycle}} = \frac{Tn}{m_{\text{cycle}}}, \quad \text{and} \quad T_{\text{dist}} = \frac{T}{m_{\text{dist}}}. \quad (16)$$

Plugging the above iteration counts into the earlier bound (8) and Corollaries 2 and 3 via the sharper result (15), we can provide upper bounds (to constant factors) on the expected optimization accuracy after T units of time for each of the distributed architectures as in Table 1. Asymptotically in the number of units of time T , both the cyclic and locally communicating stochastic optimization schemes have the same convergence rate. However, topological considerations show that the locally communicating method (Figs. 2 and 3) has better performance than the cyclic architecture, though it requires more worker coordination. Since the lower order terms matter only for large n or small T , we compare the terms $n^{2/3}/T^{2/3}$ and $D^{2/3}/T^{2/3}$ for the cyclic and locally averaged algorithms, respectively. Since $D \leq n$ for any network, the locally averaged algorithm always guarantees better performance than the cyclic algorithm. For specific graph topologies, however, we can quantify the time improvements:

- n -node cycle or path: $D = n$ so that both methods have the same convergence rate.
- \sqrt{n} -by- \sqrt{n} grid: $D = \sqrt{n}$, so the distributed method has a factor of $n^{2/3}/n^{1/3} = n^{1/3}$ improvement over the cyclic architecture.
- Balanced trees and expander graphs: $D = \mathcal{O}(\log n)$, so the distributed method has a factor—ignoring logarithmic terms—of $n^{2/3}$ improvement over cyclic.

Naturally, it is possible to modify our assumptions. In a network in which communication is cheap, or conversely, in a problem for which the computation of $\nabla F(x; \xi)$ is more expensive than communication, then the number of samples $\xi \sim P$ for which each worker computes gradients is small. Such problems are frequent in statistical machine learning, such as when learning conditional random field models, which are useful in natural language processing, computational biology, and other application areas [LMP01]. In this case, it is reasonable to have $m_{\text{cycle}} =$

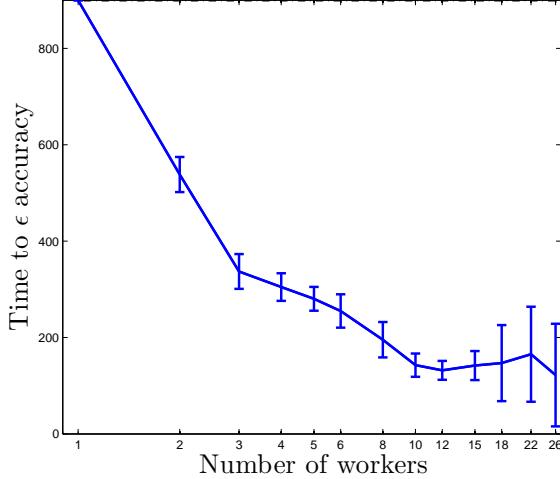


Figure 4. Optimization performance of the delayed cyclic method (9) for the Reuters RCV1 dataset when we assume that the cost of communication to the master is the same as computing the gradient of one term in the objective (17). The number of samples m computed is equal to n for each worker. Plotted is the estimated time to ϵ -accuracy as a function of number of workers n .

$\mathcal{O}(1)$, in which case $T_{\text{cycle}} = Tn$ and the cyclic delayed architecture has stronger convergence guarantees of $\mathcal{O}(\min\{n^2/T, 1/T^{2/3}\} + 1/\sqrt{Tn})$. In any case, both non-centralized protocols enjoy significant asymptotically faster convergence rates for stochastic optimization problems in spite of asynchronous delays.

5 Numerical Results

Though this paper focuses mostly on the theoretical analysis of the methods we have presented, it is important to understand the practical aspects of the above methods in solving real-world tasks and problems with real data. To that end, we use the cyclic delayed method (12) to solve a common statistical machine learning problem. Specifically, we focus on solving the logistic regression problem

$$\min_x f(x) = \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-b_i \langle a_i, x \rangle)) \quad \text{subject to } \|x\|_2 \leq R. \quad (17)$$

We use the Reuters RCV1 dataset [LYRL04], which consists of $N \approx 800000$ news articles, each labeled with some combination of the four labels economics, government, commerce, and medicine. In the above example, the vectors $a_i \in \{0, 1\}^d$, $d \approx 10^5$, are feature vectors representing the words in each article, and the labels b_i are 1 if the article is about government, -1 otherwise.

We simulate the cyclic delayed optimization algorithm (9) for the problem (17) for several choices of the number of workers n and the number of samples m computed at each worker. We summarize the results of our experiments in Figure 4. To generate the figure, we fix an ϵ (in this case, $\epsilon = .05$), then measure the time it takes the stochastic algorithm (9) to output an \hat{x} such that $f(\hat{x}) \leq \inf_{x \in \mathcal{X}} f(x) + \epsilon$. We perform each experiment ten times.

After computing the number of iterations required to achieve ϵ -accuracy, we convert the results to running time by assuming it takes one unit of time to compute the gradient of one term in the sum

defining the objective (17). We also assume that it takes 1 unit of time, i.e. $C = 1$, to communicate from one of the workers to the master, for the master to perform an update, and communicate back to one of the workers. In an n node system where each worker computes m samples of the gradient, the master receives an update every $\max\{\frac{m}{n}, 1\}$ time units. A centralized algorithm computing m samples of its gradient performs an update every m time units. By multiplying the number of iterations to ϵ -optimality by $\max\{\frac{m}{n}, 1\}$ for the distributed method and by m for the centralized, we can estimate the amount of time it takes each algorithm to achieve an ϵ -accurate solution.

We now turn to discussing Figure 4. The delayed update (9) enjoys speedup (the ratio of time to ϵ -accuracy for an n -node system versus the centralized procedure) nearly linear in the number n of worker machines until $n \geq 15$ or so. Since we use the stepsize choice $\eta(t) \propto \sqrt{t/n}$, which yields the predicted convergence rate given by Corollary 2, the $n^2m/T \approx n^3/T$ term in the convergence rate presumably becomes non-negligible for larger n . This expands on earlier experimental work with a similar method [LSZ09], which experimentally demonstrated linear speedup for small values of n , but did not investigate larger network sizes. Roughly, as predicted by our theory, for non-asymptotic regimes the cost of communication and delays due to using n nodes mitigate some of the benefits of parallelization. Nevertheless, as our analysis shows, allowing delayed and asynchronous updates still gives significant performance improvements.

6 Delayed Updates for Smooth Optimization

In this section, we prove Theorems 1 and 2. We collect in Appendix A a few technical results relevant to our proof; we will refer to results therein without comment. Before proving either theorem, we state the lemma that is the key to our argument. Lemma 4 shows that certain gradient-differencing terms are essentially of second order. As a consequence, when we combine the results of the lemma with Lemma 7, which bounds $\mathbb{E}[\|x(t) - x(t + \tau)\|^2]$, the gradient differencing terms become $\mathcal{O}(\log T)$ for step size choice $\eta(t) \propto \sqrt{t}$, or $\mathcal{O}(1)$ for $\eta(t) \equiv \eta\sqrt{T}$.

Lemma 4. *Let assumptions A and B on the function f and the compactness assumption C hold. Then for any sequence $x(t)$*

$$\sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t - \tau)), x(t + 1) - x^* \rangle \leq \frac{L}{2} \sum_{t=1}^T \|x(t - \tau) - x(t + 1)\|^2 + 2\tau GR.$$

Consequently, if $\mathbb{E}[\|x(t) - x(t + 1)\|^2] \leq \kappa(t)^2 G^2$ for a non-increasing sequence $\kappa(t)$,

$$\mathbb{E} \left[\sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t - \tau)), x(t + 1) - x^* \rangle \right] \leq \frac{LG^2(\tau + 1)^2}{2} \sum_{t=1}^T \kappa(t - \tau)^2 + 2\tau GR.$$

Proof The proof follows by using a few Bregman divergence identities to rewrite the left hand side of the above equations, then recognizing that the result is close to a telescoping sum. Recalling the definition of a Bregman divergence (2), we note the following well-known four term equality, a consequence of straightforward algebra: for any a, b, c, d ,

$$\langle \nabla f(a) - \nabla f(b), c - d \rangle = D_f(d, a) - D_f(d, b) - D_f(c, a) + D_f(c, b). \quad (18)$$

Using the equality (18), we see that

$$\begin{aligned} & \langle \nabla f(x(t)) - \nabla f(x(t-\tau)), x(t+1) - x^* \rangle \\ &= D_f(x^*, x(t)) - D_f(x^*, x(t-\tau)) - D_f(x(t+1), x(t)) + D_f(x(t+1), x(t-\tau)). \end{aligned} \quad (19)$$

To make (19) useful, we note that the Lipschitz continuity of ∇f implies

$$f(x(t+1)) \leq f(x(t-\tau)) + \langle \nabla f(x(t-\tau)), x(t+1) - x(t-\tau) \rangle + \frac{L}{2} \|x(t-\tau) - x(t+1)\|^2$$

so that recalling the definition of D_f (2) we have

$$D_f(x(t+1), x(t-\tau)) \leq \frac{L}{2} \|x(t-\tau) - x(t+1)\|^2.$$

In particular, using the non-negativity of $D_f(x, y)$, we can replace (19) with the bound

$$\langle \nabla f(x(t)) - \nabla f(x(t-\tau)), x(t+1) - x^* \rangle \leq D_f(x^*, x(t)) - D_f(x^*, x(t-\tau)) + \frac{L}{2} \|x(t-\tau) - x(t+1)\|^2.$$

Summing the inequality, we see that

$$\sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t-\tau)), x(t+1) - x^* \rangle \leq \sum_{t=T-\tau+1}^T D_f(x^*, x(t)) + \frac{L}{2} \sum_{t=1}^T \|x(t-\tau) - x(t+1)\|^2. \quad (20)$$

To bound the first Bregman divergence term, we recall that by Assumption C and the strong convexity of ψ , $\|x^* - x(t)\|^2 \leq 2D_\psi(x^*, x(t)) \leq 2R^2$, and hence the optimality of x^* implies

$$D_f(x^*, x(t)) = f(x^*) - f(x(t)) - \langle \nabla f(x(t)), x^* - x(t) \rangle \leq \|\nabla f(x(t))\|_* \|x^* - x(t)\| \leq 2GR.$$

This gives the first bound of the lemma. For the second bound, using convexity, we see that

$$\|x(t-\tau) - x(t+1)\|^2 \leq (\tau+1)^2 \sum_{s=0}^{\tau} \frac{1}{\tau+1} \|x(t-s) - x(t-s+1)\|^2,$$

so by taking expectations we have $\mathbb{E}[\|x(t) - x(t+\tau+1)\|^2] \leq (\tau+1)^2 \kappa(t-\tau)^2 G^2$. Since κ is non-increasing (by the definition of the update scheme) we see that the sum (20) is further bounded by $2\tau GR + \frac{L}{2} \sum_{t=1}^T G^2(\tau+1)^2 \kappa(t-\tau)^2$ as desired. \square

6.1 Proof of Theorem 1

The essential idea in this proof is to use convexity and smoothness to bound $f(x(t)) - f(x^*)$, then use the sequence $\{\eta(t)\}$, which decreases the stepsize $\alpha(t)$, to cancel variance terms. To begin, we define the error $e(t)$

$$e(t) := \nabla f(x(t)) - g(t-\tau)$$

where $g(t-\tau) = \nabla F(x(t-\tau); \xi(t))$ for some $\xi(t) \sim P$. Note that $e(t)$ does not have zero expectation, as there is a time delay.

By using the convexity of f and then the L -Lipschitz continuity of ∇f , for any $x^* \in \mathcal{X}$, we have

$$\begin{aligned} f(x(t)) - f(x^*) &\leq \langle \nabla f(x(t)), x(t) - x^* \rangle = \langle \nabla f(x(t)), x(t+1) - x^* \rangle + \langle \nabla f(x(t)), x(t) - x(t+1) \rangle \\ &\leq \langle \nabla f(x(t)), x(t+1) - x^* \rangle + f(x(t)) - f(x(t+1)) + \frac{L}{2} \|x(t) - x(t+1)\|^2, \end{aligned}$$

so that

$$\begin{aligned} f(x(t+1)) - f(x^*) &\leq \langle \nabla f(x(t)), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2 \\ &= \langle g(t-\tau), x(t+1) - x^* \rangle + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2 \\ &= \langle z(t+1), x(t+1) - x^* \rangle - \langle z(t), x(t+1) - x^* \rangle + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2. \end{aligned}$$

Now, by applying Lemma 5 in Appendix A and the definition of the update (9), we see that

$$-\langle z(t), x(t+1) - x^* \rangle \leq -\langle z(t), x(t) - x^* \rangle + \frac{1}{\alpha(t)} [\psi(x(t+1)) - \psi(x(t))] - \frac{1}{\alpha(t)} D_\psi(x(t+1), x(t)),$$

which implies

$$\begin{aligned} f(x(t+1)) - f(x^*) &\leq \langle z(t+1), x(t+1) - x^* \rangle - \langle z(t), x(t) - x^* \rangle + \frac{1}{\alpha(t)} [\psi(x(t+1)) - \psi(x(t))] \\ &\quad - L D_\psi(x(t+1), x(t)) - \eta(t) D_\psi(x(t+1), x(t)) + \frac{L}{2} \|x(t) - x(t+1)\|^2 + \langle e(t), x(t+1) - x^* \rangle \\ &\leq \langle z(t+1), x(t+1) - x^* \rangle - \langle z(t), x(t) - x^* \rangle + \frac{1}{\alpha(t)} [\psi(x(t+1)) - \psi(x(t))] \\ &\quad - \eta(t) D_\psi(x(t+1), x(t)) + \langle e(t), x(t+1) - x^* \rangle. \end{aligned} \tag{21}$$

To get the bound (21), we substituted $\alpha(t)^{-1} = L + \eta(t)$ and then used the fact that ψ is strongly convex, so $D_\psi(x(t+1), x(t)) \geq \frac{1}{2} \|x(t) - x(t+1)\|^2$. By summing the bound (21), we have the following non-probabilistic inequality:

$$\begin{aligned} &\sum_{t=1}^T f(x(t+1)) - f(x^*) \\ &\leq \langle z(T+1), x(T+1) - x^* \rangle + \frac{1}{\alpha(T)} \psi(x(T+1)) + \sum_{t=1}^T \psi(x(t)) \left[\frac{1}{\alpha(t-1)} - \frac{1}{\alpha(t)} \right] \\ &\quad - \sum_{t=1}^T \eta(t) D_\psi(x(t+1), x(t)) + \sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle \\ &\leq \frac{1}{\alpha(T+1)} \psi(x^*) + \sum_{t=1}^T \psi(x(t)) \left[\frac{1}{\alpha(t-1)} - \frac{1}{\alpha(t)} \right] - \sum_{t=1}^T \eta(t) D_\psi(x(t+1), x(t)) \\ &\quad + \sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle \end{aligned} \tag{22}$$

since $\psi(x) \geq 0$ and $x(T+1)$ minimizes $\langle z(T+1), x \rangle + \frac{1}{\alpha(T+1)}\psi(x)$. What remains is to control the summed $e(t)$ terms in the bound (22). We can do this simply using the second part of Lemma 4. Indeed, we have

$$\begin{aligned} & \sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle \\ &= \sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t-\tau)), x(t+1) - x^* \rangle + \sum_{t=1}^T \langle \nabla f(x(t-\tau)) - g(t-\tau), x(t+1) - x^* \rangle. \end{aligned} \quad (23)$$

We can apply Lemma 4 to the first term in (23) by bounding $\|x(t) - x(t+1)\|$ with Lemma 7. Since $\eta(t) \propto \sqrt{t+\tau}$, Lemma 7 with $t_0 = \tau$ implies $\mathbb{E}[\|x(t) - x(t+1)\|^2] \leq \frac{4G^2}{\eta(t)^2}$. As a consequence,

$$\mathbb{E} \left[\sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t-\tau)), x(t+1) - x^* \rangle \right] \leq 2\tau GR + 2L(\tau+1)^2 G^2 \sum_{t=1}^T \frac{1}{\eta(t-\tau)^2}.$$

What remains, then, is to bound the stochastic (second) term in (23). This is straightforward, though:

$$\begin{aligned} & \langle \nabla f(x(t-\tau)) - g(t-\tau), x(t+1) - x^* \rangle \\ &= \langle \nabla f(x(t-\tau)) - g(t-\tau), x(t) - x^* \rangle + \langle \nabla f(x(t-\tau)) - g(t-\tau), x(t+1) - x(t) \rangle \\ &\leq \langle \nabla f(x(t-\tau)) - g(t-\tau), x(t) - x^* \rangle + \frac{1}{2\eta(t)} \|\nabla f(x(t-\tau)) - g(t-\tau)\|_*^2 + \frac{\eta(t)}{2} \|x(t+1) - x(t)\|^2 \end{aligned}$$

by the Fenchel-Young inequality applied to the conjugate pair $\frac{1}{2}\|\cdot\|_*^2$ and $\frac{1}{2}\|\cdot\|^2$. In addition, $\nabla f(x(t-\tau)) - g(t-\tau)$ is independent of $x(t)$ given the sigma-field containing $g(1), \dots, g(t-\tau-1)$, since $x(t)$ is a function of gradients to time $t-\tau-1$, so the first term has zero expectation. Also recall that $\mathbb{E}[\|\nabla f(x(t-\tau)) - g(t-\tau)\|_*^2]$ is bounded by σ^2 by assumption. Combining the above two bounds into (23), we see that

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}[\langle e(t), x(t+1) - x^* \rangle] \\ &\leq \frac{\sigma^2}{2} \sum_{t=1}^T \frac{1}{\eta(t)} + \frac{1}{2} \sum_{t=1}^T \eta(t) \|x(t+1) - x(t)\|^2 + 2LG^2(\tau+1)^2 \sum_{t=1}^T \frac{1}{\eta(t-\tau)^2} + 2\tau GR. \end{aligned} \quad (24)$$

Since $D_\psi(x(t+1), x(t)) \geq \frac{1}{2}\|x(t) - x(t+1)\|^2$, combining (24) with (22) and noting that $\frac{1}{\alpha(t-1)} - \frac{1}{\alpha(t)} \leq 0$ gives

$$\sum_{t=1}^T \mathbb{E}f(x(t+1)) - f(x^*) \leq \frac{1}{\alpha(T+1)}\psi(x^*) + \frac{\sigma^2}{2} \sum_{t=1}^T \frac{1}{\eta(t)} + 2LG^2(\tau+1)^2 \sum_{t=1}^T \frac{1}{\eta(t-\tau)^2} + 2\tau GR.$$

6.2 Proof of Theorem 2

The proof of Theorem 2 is similar to that of Theorem 1, so we will be somewhat terse. We define the error $e(t) = \nabla f(x(t)) - g(t-\tau)$, identically as in the earlier proof, and begin as we did in the

proof of Theorem 1. Recall that

$$f(x(t+1)) - f(x^*) \leq \langle g(t-\tau), x(t+1) - x^* \rangle + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2. \quad (25)$$

Applying the first-order optimality condition to the definition of $x(t+1)$ (5), we get

$$\langle \alpha(t)g(t-\tau) + \nabla\psi(x(t+1)) - \nabla\psi(x(t)), x - x(t+1) \rangle \geq 0$$

for all $x \in \mathcal{X}$. In particular, we have

$$\begin{aligned} \alpha(t) \langle g(t-\tau), x(t+1) - x^* \rangle &\leq \langle \nabla\psi(x(t+1)) - \nabla\psi(x(t)), x^* - x(t+1) \rangle \\ &= D_\psi(x^*, x(t)) - D_\psi(x^*, x(t+1)) - D_\psi(x(t+1), x(t)). \end{aligned}$$

Applying the above to the inequality (25), we see that

$$\begin{aligned} f(x(t+1)) - f(x^*) &\leq \frac{1}{\alpha(t)} [D_\psi(x^*, x(t)) - D_\psi(x^*, x(t+1)) - D_\psi(x(t+1), x(t))] + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2 \\ &\leq \frac{1}{\alpha(t)} [D_\psi(x^*, x(t)) - D_\psi(x^*, x(t+1))] + \langle e(t), x(t+1) - x^* \rangle - \eta(t) D_\psi(x(t+1), x(t)) \end{aligned} \quad (26)$$

where for the last inequality, we use the fact that $D_\psi(x(t+1), x(t)) \geq \frac{1}{2} \|x(t) - x(t+1)\|^2$, by the strong convexity of ψ , and that $\alpha(t)^{-1} = L + \eta(t)$. By summing the inequality (26), we have

$$\begin{aligned} \sum_{t=1}^T f(x(t+1)) - f(x^*) &\leq \frac{1}{\alpha(1)} D_\psi(x^*, x(1)) + \sum_{t=2}^T D_\psi(x^*, x(t)) \left[\frac{1}{\alpha(t)} - \frac{1}{\alpha(t-1)} \right] \\ &\quad - \sum_{t=1}^T \eta(t) D_\psi(x(t+1), x(t)) + \sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle. \end{aligned} \quad (27)$$

Comparing the bound (27) with the earlier bound for the dual averaging algorithms (22), we see that the only essential difference is the $\alpha(t)^{-1} - \alpha(t-1)^{-1}$ terms. The compactness assumption guarantees that $D_\psi(x^*, x(t)) \leq R^2$, however, so

$$\sum_{t=2}^T D_\psi(x^*, x(t)) \left[\frac{1}{\alpha(t)} - \frac{1}{\alpha(t-1)} \right] \leq \frac{R^2}{\alpha(T)}.$$

The remainder of the proof uses Lemmas 7 and 4 completely identically to the proof of Theorem 1.

6.3 Proof of Corollary 1

We prove this result only for the mirror descent algorithm (10), as the proof for the dual-averaging-based algorithm (9) is similar. We define the error at time t to be $e(t) = \nabla f(x(t)) - g(t-\tau(t))$, and observe that we only need to control the second term involving $e(t)$ in the bound (26) differently.

Expanding the error terms above and using Fenchel's inequality as in the proofs of Theorems 1 and 2, we have

$$\begin{aligned} & \langle e(t), x(t+1) - x^* \rangle \\ & \leq \langle \nabla f(x(t)) - \nabla f(x(t - \tau(t))), x(t+1) - x^* \rangle + \langle \nabla f(x(t - \tau(t))) - g(t - \tau(t)), x(t) - x^* \rangle \\ & \quad + \frac{1}{2\eta(t)} \|\nabla f(x(t - \tau(t))) - g(t - \tau(t))\|_*^2 + \frac{\eta(t)}{2} \|x(t+1) - x(t)\|^2, \end{aligned}$$

Now we note that conditioned on the delay $\tau(t)$, we have

$$\mathbb{E}[\|x(t - \tau(t)) - x(t+1)\|^2 \mid \tau(t)] \leq G^2(\tau(t) + 1)^2 \alpha(t - \tau(t))^2.$$

Consequently we apply Lemma 4 (specifically, following the bounds (19) and (20)) and find

$$\begin{aligned} & \sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t - \tau(t))), x(t+1) - x^* \rangle \\ & \leq \sum_{t=1}^T [D_f(x^*, x(t)) - D_f(x^*, x(t - \tau(t)))] + G^2 \sum_{t=1}^T (\tau(t) + 1)^2 \alpha(t - \tau(t))^2. \end{aligned}$$

The sum of D_f terms telescopes, leaving only terms not received by the gradient procedure within T iterations, and we can use $\alpha(t) \leq \frac{1}{\eta\sqrt{T}}$ for all t to derive the further bound

$$\sum_{t:t+\tau(t)>T} D_f(x^*, x(t)) + \frac{G^2}{\eta^2 T} \sum_{t=1}^T (\tau(t) + 1)^2. \quad (28)$$

To control the quantity (28), all we need is to bound the expected cardinality of the set $\{t \in [T] : t + \tau(t) > T\}$. Using Chebyshev's inequality and standard expectation bounds, we have

$$\mathbb{E}[\text{card}(\{t \in [T] : t + \tau(t) > T\})] = \sum_{t=1}^T \mathbb{P}(t + \tau(t) > T) \leq 1 + \sum_{t=1}^{T-1} \frac{\mathbb{E}[\tau(t)^2]}{(T-t)^2} \leq 1 + 2B^2,$$

where the last inequality comes from our assumption that $\mathbb{E}[\tau(t)^2] \leq B^2$. As in Lemma 4, we have $D_f(x^*, x(t)) \leq 2GR$, which yields

$$\mathbb{E} \left[\sum_{t=1}^T \langle \nabla f(x(t)) - \nabla f(x(t - \tau(t))), x(t+1) - x^* \rangle \right] \leq 6GRB^2 + \frac{G^2(B+1)^2}{\eta^2}$$

We can control the remaining terms as in the proofs of Theorems 1 and 2.

7 Proof of Theorem 3

The proof of Theorem 3 is not too difficult given our previous work—all we need to do is redefine the error $e(t)$ and use $\eta(t)$ to control the variance terms that arise. To that end, we define the gradient error terms that we must control. In this proof, we set

$$e(t) := \nabla f(x(t)) - \sum_{i=1}^n \lambda_i g_i(t - \tau(i)) \quad (29)$$

where $g_i(t) = \nabla f(x(t); \xi_i(t))$ is the gradient of node i computed at the parameter $x(t)$ and $\tau(i)$ is the delay associated with node i .

Using Assumption B as in the proofs of previous theorems, then applying Lemma 5, we have

$$\begin{aligned}
f(x(t+1)) - f(x^*) &\leq \langle \nabla f(x(t)), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2 \\
&= \left\langle \sum_{i=1}^n \lambda_i g_i(t - \tau(i)), x(t+1) - x^* \right\rangle + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2 \\
&= \langle z(t+1), x(t+1) - x^* \rangle - \langle z(t), x(t+1) - x^* \rangle + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2 \\
&\leq \langle z(t+1), x(t+1) - x^* \rangle - \langle z(t), x(t) - x^* \rangle + \frac{1}{\alpha(t)} \psi(x(t+1)) - \frac{1}{\alpha(t)} \psi(x(t)) \\
&\quad - \frac{1}{\alpha(t)} D_\psi(x(t+1), x(t)) + \langle e(t), x(t+1) - x^* \rangle + \frac{L}{2} \|x(t) - x(t+1)\|^2.
\end{aligned}$$

We telescope as in the proofs of Theorems 1 and 2, canceling $\frac{L}{2} \|x(t) - x(t+1)\|^2$ with the LD_ψ divergence terms to see that

$$\begin{aligned}
&\sum_{t=1}^T f(x(t+1)) - f(x^*) \\
&\leq \langle z(T+1), x(T+1) - x^* \rangle + \frac{1}{\alpha(T)} \psi(x(T)) - \sum_{t=1}^T \eta(t) D_\psi(x(t+1), x(t)) + \sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle \\
&\leq \frac{1}{\alpha(T+1)} \psi(x^*) - \sum_{t=1}^T \eta(t) D_\psi(x(t+1), x(t)) + \sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle. \tag{30}
\end{aligned}$$

This is exactly as in the non-probabilistic bound (22) from the proof of Theorem 1, but the definition (29) of the error $e(t)$ here is different.

What remains is to control the error term in (30). Writing the terms out, we have

$$\begin{aligned}
\sum_{t=1}^T \langle e(t), x(t+1) - x^* \rangle &= \sum_{t=1}^T \left\langle \nabla f(x(t)) - \sum_{i=1}^n \lambda_i \nabla f(x(t - \tau(i))), x(t+1) - x^* \right\rangle \\
&\quad + \sum_{t=1}^T \left\langle \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))], x(t+1) - x^* \right\rangle \tag{31}
\end{aligned}$$

Bounding the first term above is simple via Lemma 4: as in the proof of Theorem 1 earlier, we have

$$\begin{aligned}
&\mathbb{E} \left[\sum_{t=1}^T \left\langle \nabla f(x(t)) - \sum_{i=1}^n \lambda_i \nabla f(x(t - \tau(i))), x(t+1) - x^* \right\rangle \right] \\
&= \sum_{i=1}^n \lambda_i \sum_{t=1}^T \mathbb{E} [\langle \nabla f(x(t)) - \nabla f(x(t - \tau(i))), x(t+1) - x^* \rangle] \\
&\leq 2 \sum_{i=1}^n \lambda_i L G^2 (\tau(i) + 1)^2 \sum_{t=1}^T \frac{1}{\eta(t - \tau)^2} + \sum_{i=1}^n \lambda_i 2\tau(i) G R.
\end{aligned}$$

We use the same technique as the proof of Theorem 1 to bound the second term from (31). Indeed, the Fenchel-Young inequality gives

$$\begin{aligned}
& \left\langle \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))], x(t+1) - x^* \right\rangle \\
&= \left\langle \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))], x(t) - x^* \right\rangle \\
&\quad + \left\langle \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))], x(t+1) - x(t) \right\rangle \\
&\leq \left\langle \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))], x(t) - x^* \right\rangle \\
&\quad + \frac{1}{2\eta(t)} \left\| \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))] \right\|_*^2 + \frac{\eta(t)}{2} \|x(t+1) - x(t)\|^2.
\end{aligned}$$

By assumption, given the information at worker i at time $t - \tau(i)$, $g_i(t - \tau(i))$ is independent of $x(t)$, so the first term has zero expectation. More formally, this happens because $x(t)$ is a function of gradients $g_i(1), \dots, g_i(t - \tau(i) - 1)$ from each of the nodes i and hence the expectation of the first term conditioned on $\{g_i(1), \dots, g_i(t - \tau(i) - 1)\}_{i=1}^n$ is 0. The last term is canceled by the Bregman divergence terms in (30), so combining the bound (31) with the above two paragraphs yields

$$\begin{aligned}
\sum_{t=1}^T \mathbb{E} f(x(t+1)) - f(x^*) &\leq \frac{1}{\alpha(t)} \psi(x^*) + 2 \sum_{i=1}^n \lambda_i L G^2 (\tau(i) + 1)^2 \sum_{t=1}^T \frac{1}{\eta(t - \tau)^2} + \sum_{i=1}^n \lambda_i 2\tau(i) G R \\
&\quad + \sum_{t=1}^T \frac{1}{2\eta(t)} \mathbb{E} \left\| \sum_{i=1}^n \lambda_i [\nabla f(x(t - \tau(i))) - g_i(t - \tau(i))] \right\|_*^2.
\end{aligned}$$

8 Conclusion and Discussion

In this paper, we have studied dual averaging and mirror descent algorithms for smooth and non-smooth stochastic optimization in delayed settings, showing applications of our results to distributed optimization. We showed that for smooth problems, we can preserve the performance benefits of parallelization over centralized stochastic optimization even when we relax synchronization requirements. Specifically, we presented methods that take advantage of distributed computational resources and are robust to node failures, communication latency, and node slowdowns. In addition, by distributing computation for stochastic optimization problems, we were able to exploit asynchronous processing without incurring any asymptotic penalty due to the delays incurred. In addition, though we omit these results for brevity, it is possible to extend all of our expected convergence results to guarantees with high-probability.

Acknowledgments

In performing this research, AA was supported by a Microsoft Research Fellowship, and JCD was supported by the National Defense Science and Engineering Graduate Fellowship (NDSEG)

Program. We are very grateful to Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao for illuminating conversations on distributed stochastic optimization and communication of their proof of the bound (8). We would also like to thank Yoram Singer for reading a draft of this manuscript and giving useful feedback.

A Technical Results about Proximal Functions

In this section, we collect several useful results about proximal functions and continuity properties of the solutions of proximal operators. We give proofs of all uncited results in Appendix B. We begin with results useful for the dual-averaging updates (4) and (9).

We define the proximal dual function

$$\psi_\alpha^*(z) := \sup_{x \in \mathcal{X}} \left\{ \langle -z, x \rangle - \frac{1}{\alpha} \psi(x) \right\}. \quad (32)$$

Since $\nabla \psi_\alpha^*(z) = \operatorname{argmax}_{x \in \mathcal{X}} \{ \langle -z, x \rangle - \alpha^{-1} \psi(x) \}$, it is clear that $x(t) = \nabla \psi_{\alpha(t)}^*(z(t))$. Further by strong convexity of ψ , we have that $\nabla \psi_\alpha^*(z)$ is α -Lipschitz continuous [Nes09, HUL96b, Chapter X], that is, for the norm $\|\cdot\|$ with respect to which ψ is strongly convex and its associated dual norm $\|\cdot\|_*$,

$$\|\nabla \psi_\alpha^*(y) - \nabla \psi_\alpha^*(z)\| \leq \alpha \|y - z\|_*. \quad (33)$$

We will find one more result about solutions to the dual averaging update useful. This result has essentially been proven in many contexts [Nes09, Tse08, DGBSX10a].

Lemma 5. *Let x^+ minimize $\langle z, x \rangle + A\psi(x)$ for all $x \in \mathcal{X}$. Then for any $x \in \mathcal{X}$,*

$$\langle z, x \rangle + A\psi(x) \geq \langle z, x^+ \rangle + A\psi(x^+) + AD_\psi(x, x^+)$$

Now we turn to describing properties of the mirror-descent step (5), which we will also use frequently. The lemma allows us to bound differences between $x(t)$ and $x(t+1)$ for the mirror-descent family of algorithms.

Lemma 6. *Let x^+ minimize $\langle g, x \rangle + \frac{1}{\alpha} D_\psi(x, y)$ over $x \in \mathcal{X}$. Then $\|x^+ - y\| \leq \alpha \|g\|_*$.*

The last technical lemma we give explicitly bounds the differences between $x(t)$ and $x(t+\tau)$, for some $\tau \geq 1$, by using the above continuity lemmas.

Lemma 7. *Let Assumption A hold. Define $x(t)$ via the dual-averaging updates (4), (9), or (12) or the mirror-descent updates (5), (10), or (13). Let $\alpha(t)^{-1} = L + \eta(t+t_0)^c$ for some $c \in [0, 1]$, $\eta > 0$, $t_0 \geq 0$, and $L \geq 0$. Then for any fixed τ ,*

$$\mathbb{E}[\|x(t) - x(t+\tau)\|^2] \leq \frac{4G^2\tau^2}{\eta^2(t+t_0)^{2c}} \quad \text{and} \quad \mathbb{E}[\|x(t) - x(t+\tau)\|] \leq \frac{2G\tau}{\eta(t+t_0)^c}.$$

B Proofs of Proximal Operator Properties

Proof of Lemma 6 The inequality is clear when $x^+ = y$, so assume that $x^+ \neq y$. Since x^+ minimizes $\langle g, x \rangle + \frac{1}{\alpha} D_\psi(x, y)$, the first order conditions for optimality imply

$$\langle \alpha g + \nabla \psi(x^+) - \nabla \psi(y), x - x^+ \rangle \geq 0$$

for any $x \in \mathcal{X}$. Thus we can choose $y = x$ and see that

$$\alpha \langle g, y - x \rangle \geq \langle \nabla \psi(x^+) - \nabla \psi(y), x^+ - y \rangle \geq \|x^+ - y\|^2,$$

where the last inequality follows from the strong convexity of ψ . Using Hölder's inequality gives that $\alpha \|g\|_* \|y - x\| \geq \|x^+ - y\|^2$, and dividing by $\|y - x\|$ completes the proof. \square

Proof of Lemma 7 We first show the lemma for the dual-averaging updates. Recall that $x(t) = \nabla \psi_{\alpha(t)}^*(z(t))$ and $\nabla \psi_\alpha^*$ is α -Lipschitz continuous. Using the triangle inequality,

$$\begin{aligned} \|x(t) - x(t + \tau)\| &= \left\| \nabla \psi_{\alpha(t)}^*(z(t)) - \nabla \psi_{\alpha(t+\tau)}^*(z(t + \tau)) \right\| \\ &= \left\| \nabla \psi_{\alpha(t)}^*(z(t)) - \nabla \psi_{\alpha(t+\tau)}^*(z(t)) + \nabla \psi_{\alpha(t+\tau)}^*(z(t)) - \nabla \psi_{\alpha(t+\tau)}^*(z(t + \tau)) \right\| \\ &\leq \left\| \nabla \psi_{\alpha(t)}^*(z(t)) - \nabla \psi_{\alpha(t+\tau)}^*(z(t)) \right\| + \left\| \nabla \psi_{\alpha(t+\tau)}^*(z(t)) - \nabla \psi_{\alpha(t+\tau)}^*(z(t + \tau)) \right\| \\ &\leq (\alpha(t) - \alpha(t + \tau)) \|z(t)\|_* + \alpha(t + \tau) \|z(t) - z(t + \tau)\|_*. \end{aligned} \quad (34)$$

It is easy to check that for $c \in [0, 1]$,

$$\alpha(t) - \alpha(t + \tau) \leq \frac{c\eta\tau}{(L + \eta t^c)^2 t^{1-c}} \leq \frac{c\tau}{\eta t^{1+c}}.$$

By convexity of $\|\cdot\|_*^2$, we can bound $\mathbb{E}[\|z(t) - z(t + \tau)\|_*^2]$:

$$\mathbb{E}[\|z(t) - z(t + \tau)\|_*^2] = \tau^2 \mathbb{E} \left[\left\| \frac{1}{\tau} \sum_{s=1}^{\tau} z(t+s) - z(t+s-1) \right\|_*^2 \right] = \tau^2 \mathbb{E} \left[\left\| \frac{1}{\tau} \sum_{s=0}^{\tau-1} g(s) \right\|_*^2 \right] \leq \tau^2 G^2,$$

since $\mathbb{E}[\|\partial F(x; \xi)\|_*^2] \leq G^2$ by assumption. Thus, bound (34) gives

$$\begin{aligned} \mathbb{E}[\|x(t) - x(t + \tau)\|^2] &\leq 2(\alpha(t) - \alpha(t + \tau))^2 \mathbb{E}[\|z(t)\|_*^2] + 2\alpha(t + \tau)^2 \mathbb{E}[\|z(t) - z(t + \tau)\|_*^2] \\ &\leq \frac{2c^2 t^2 \tau^2 G^2}{\eta^2 t^{2+2c}} + 2G^2 \tau^2 \alpha(t + \tau)^2 = \frac{2c^2 \tau^2 G^2}{\eta^2 t^{2c}} + \frac{2G^2 \tau^2}{(L + \eta(t + \tau)^c)^2}, \end{aligned}$$

where we use Cauchy-Schwarz inequality in the first step. Since $c \leq 1$, the last term is clearly bounded by $4G^2 \tau^2 / \eta^2 t^{2c}$.

To get the slightly tighter bound on the first moment in the statement of the lemma, simply use the triangle inequality from the bound (34) and that $\sqrt{\mathbb{E}X^2} \geq \mathbb{E}|X|$.

The proof for the mirror-descent family of updates is similar. We focus on non-delayed update (5), as the other updates simply modify the indexing of $g(t+s)$ below. We know from Lemma 6 and the triangle inequality that

$$\|x(t) - x(t + \tau)\| \leq \sum_{s=1}^{\tau} \|x(t+s) - x(t+s-1)\| \leq \sum_{s=1}^{\tau} \alpha(t+s-1) \|g(t+s)\|_*$$

Squaring the above bound, taking expectations, and recalling that $\alpha(t)$ is non-increasing, we see

$$\begin{aligned} \mathbb{E}[\|x(t) - x(t + \tau)\|^2] &\leq \sum_{s=1}^{\tau} \sum_{r=1}^{\tau} \alpha(t+s)\alpha(t+r) \mathbb{E}[\|g(t+s)\|_* \|g(t+r)\|_*] \\ &\leq \tau^2 \alpha(t)^2 \max_{r,s} \sqrt{\mathbb{E}[\|g(t+s)\|_*^2]} \sqrt{\mathbb{E}[\|g(t+r)\|_*^2]} \leq \tau^2 \alpha(t)^2 G^2 \end{aligned}$$

by Hölder's inequality. Substituting the appropriate value for $\alpha(t)$ completes the proof. \square

C Error in [LSZ09]

Langford et al. [LSZ09], in Lemma 1 of their paper, state an upper bound on $\langle g(t - \tau), x(t - \tau) - x^* \rangle$ that is essential to the proofs of all of their results. However, the lemma only holds as an equality for unconstrained optimization (i.e. when the set $\mathcal{X} = \mathbb{R}^d$); in the presence of constraints, it fails to hold (even as an upper bound). To see why, we consider a simple one-dimensional example with $\mathcal{X} = [-1, 1]$, $f(x) = |x|$, $\eta \equiv 2$ and we evaluate both sides of their lemma with $\tau = 1$ and $t = 2$. The left hand side of their bound evaluates to 1, while the right hand side is -1 , and the inequality claimed in the lemma fails. The proofs of their main theorems rely on the application of their Lemma 1 with equality, restricting those results only to unconstrained optimization. However, the results also require boundedness of the gradients $g(t)$ over all of \mathcal{X} as well as boundedness of the distance between the iterates $x(t)$. Few convex functions have bounded gradients over all of \mathbb{R}^d ; and without constraints the iterates $x(t)$ are seldom bounded for all iterations t .

References

- [ABRW10] A. Agarwal, P. Bartlett, P. Ravikumar, and M. Wainwright, *Information-theoretic lower bounds on the oracle complexity of convex optimization*, Submitted to *IEEE Transactions on Information Theory*, URL <http://arxiv.org/abs/1009.0571>, 2010.
- [Ber73] D. P. Bertsekas, *Stochastic optimization problems with nondifferentiable cost functionals*, Journal of Optimization Theory and Applications **12** (1973), no. 2, 218–231.
- [Bre67] L. M. Bregman, *The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming*, USSR Computational Mathematics and Mathematical Physics **7** (1967), 200–217.
- [BT89] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*, Prentice-Hall, Inc., 1989.
- [BT03] A. Beck and M. Teboulle, *Mirror descent and nonlinear projected subgradient methods for convex optimization*, Operations Research Letters **31** (2003), 167–175.
- [DAW10] J. Duchi, A. Agarwal, and M. Wainwright, *Dual averaging for distributed optimization: convergence analysis and network scaling*, URL <http://arxiv.org/abs/1005.2012>, 2010.
- [DGBSX10a] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, *Optimal distributed online prediction using mini-batches*, URL <http://arxiv.org/abs/1012.1367>, 2010.
- [DGBSX10b] _____, *Robust distributed online prediction*, URL <http://arxiv.org/abs/1012.1370>, 2010.
- [HTF01] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning*, Springer, 2001.

[HUL96a] J. Hiriart-Urruty and C. Lemaréchal, *Convex Analysis and Minimization Algorithms I*, Springer, 1996.

[HUL96b] ———, *Convex Analysis and Minimization Algorithms II*, Springer, 1996.

[JNT08] A. Juditsky, A. Nemirovski, and C. Tauvel, *Solving variational inequalities with the stochastic mirror-prox algorithm*, URL <http://arxiv.org/abs/0809.0815>, 2008.

[JRJ09] B. Johansson, M. Rabi, and M. Johansson, *A randomized incremental subgradient method for distributed optimization in networked systems*, SIAM Journal on Optimization **20** (2009), no. 3, 1157–1170.

[Lan10] G. Lan, *An optimal method for stochastic composite optimization*, Mathematical Programming Series A (2010), Online first, to appear. URL http://www.ise.ufl.edu/glan/papers/OPT_SA4.pdf.

[LMP01] J. Lafferty, A. McCallum, and F. Pereira, *Conditional random fields: Probabilistic models for segmenting and labeling sequence data*, Proceedings of the Eighteenth International Conference on Machine Learning, 2001, pp. 282–289.

[LSZ09] J. Langford, A. Smola, and M. Zinkevich, *Slow learners are fast*, Advances in Neural Information Processing Systems 22, 2009, pp. 2331–2339.

[LYRL04] D. Lewis, Y. Yang, T. Rose, and F. Li, *RCV1: A new benchmark collection for text categorization research*, Journal of Machine Learning Research **5** (2004), 361–397.

[NBB01] A. Nedić, D.P. Bertsekas, and V.S. Borkar, *Distributed asynchronous incremental subgradient methods*, Inherently Parallel Algorithms in Feasibility and Optimization and their Applications (D. Butnariu, Y. Censor, and S. Reich, eds.), Studies in Computational Mathematics, vol. 8, Elsevier, 2001, pp. 381–407.

[Nes09] Y. Nesterov, *Primal-dual subgradient methods for convex problems*, Mathematical Programming A **120** (2009), no. 1, 261–283.

[NJLS09] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, *Robust stochastic approximation approach to stochastic programming*, SIAM Journal on Optimization **19** (2009), no. 4, 1574–1609.

[NO09] A. Nedić and A. Ozdaglar, *Distributed subgradient methods for multi-agent optimization*, IEEE Transactions on Automatic Control **54** (2009), 48–61.

[NY83] A. Nemirovski and D. Yudin, *Problem complexity and method efficiency in optimization*, Wiley, New York, 1983.

[Pol87] B. T. Polyak, *Introduction to optimization*, Optimization Software, Inc., 1987.

[RM51] H. Robbins and S. Monro, *A stochastic approximation method*, Annals of Mathematical Statistics **22** (1951), 400–407.

[RNV10] S. S. Ram, A. Nedić, and V. V. Veeravalli, *Distributed stochastic subgradient projection algorithms for convex optimization*, Journal of Optimization Theory and Applications **147** (2010), no. 3, 516–545.

- [Tse08] P. Tseng, *On accelerated proximal gradient methods for convex-concave optimization*, Tech. report, Department of Mathematics, University of Washington, 2008.
- [Tsi84] J. Tsitsiklis, *Problems in decentralized decision making and computation*, Ph.D. thesis, Massachusetts Institute of Technology, 1984.
- [Xia10] L. Xiao, *Dual averaging methods for regularized stochastic learning and online optimization*, Journal of Machine Learning Research **11** (2010), 2543–2596.